

A Practical Approach to the Vexing Issue of “Risk Appetite”

By Tom Coyne, Sachuest Advisors

June 2013

This essay is about a very important (but too seldom discussed) enterprise risk management issue that I have encountered as the CEO and CFO of non-financial companies: Just how can a director or executive practically apply the concept of organizational “risk appetite” that figures so prominently in many well known ERM frameworks?

I will begin by addressing some basic linguistic and intellectual confusion that underlies this issue, and then move on to some critical underlying considerations at the individual and group level. I will conclude with some practical recommendations for company directors.

Two Basic Sources of Confusion that Complicate This Discussion

I’ve always found the use of the term “risk appetite” in various enterprise risk management (ERM) frameworks to be unusual and confusing. Research in a wide variety of subjects (e.g., consumer and investor behavior, psychology, and corporate finance) often refers not to “risk appetite”, but rather to varying formulas for and degrees of risk aversion, risk seeking, and risk tolerance.

Even more important, however, is the frequent intellectual confusion in ERM discussions over the critical distinction between uncertainty and risk. In point of fact, the latter is a very small subset of the former, and refers to those situations in which a decision maker can work out the full range of possible outcomes in a situation, as well as their consequences and probabilities, with the latter summing to one. If all three of these conditions do not prevail (which is generally the case in

the real world), then we are dealing with uncertainty, not risk. This distinction was most famously made by Frank Knight in his 1921 book, Risk, Uncertainty, and Profit. It was also a focus of John Maynard Keynes' work, especially Chapter Twelve of his best known work, The General Theory of Employment, Interest, and Money. As he noted, "the outstanding fact is the extreme precariousness of the basis of the knowledge on which our estimates of prospective yield [the future profitability of an investment] have to be made. Our knowledge of the factors which will govern the yield of an investment some years hence is usually very slight, and often negligible." In my view, the best practical definition of uncertainty comes from Lipshitz and Strauss, who describe it as "a sense of doubt that threatens to block or delay action" (see "Coping With Uncertainty: A Naturalistic Decision-Making Analysis").

In recent years, researchers have identified three broad types of uncertainty. The simplest to understand is the irreducible variation that is a normal part of the operation of many physical and biological systems. To use a medical example, a person who takes their blood pressure at the same time every day will obtain a range of measurements, even if nothing else in their life (e.g., weight, stress, etc.) has changed. Closely related to this is the variability that is inherent in any statistical sampling process (e.g., this is what produces the margin of error reported in political polls).

The second source of uncertainty is due to our lack of knowledge about a process or system. A typical example of this is our uncertainty about how to describe the range of possible outcomes for a variable in a model (e.g., assumptions about the form of the probability distribution to use, and the values to include in it). For most people, this is the type of uncertainty that most readily comes to mind.

The third source of uncertainty is the extent to which our explicit or implicit model of a given problem or system corresponds to its true nature – i.e., do we really understand what is going on? This type of uncertainty is the realm of Donald

Rumsfeld's famous "unknown unknowns", and it is undoubtedly the most psychologically unsettling for most people.

In the case of stable physical systems, this model uncertainty, like parameter uncertainty, can be reduced over time through the collection of additional data. However, this is not the case for systems that are constantly evolving, like technological, social, political, business and economic systems, in which accurate prediction is extremely difficult. All of these are so called "complex adaptive systems", which are populated by intelligent agents who change their goals and strategies over time, and which are characterized by dense networks of cause and effect relationships, many of which are non-linear and time delayed in their impact. The behavior of such systems is said to be "emergent", because at the aggregate/macro level, it cannot be predicted from micro-level knowledge of individual agents' decision rules. It is also said to be emergent because the nature of the system itself constantly evolves over time. Moreover, many complex adaptive systems produce patterns of outcomes that are characterized not by the familiar normal/Gaussian/bell curve distribution, but rather by exponential/Pareto/power law distributions, which have far more "tail risk" (i.e., potentially extreme outcomes) than we intuitively expect. More generally, research has shown that human beings have a very difficult time explaining and predicting the operation of complex adaptive systems (e.g., see "Complex Dynamics in Learning Complicated Games" by Galla and Farmer, "Learning from Evidence in a Complex World" by John Sterman, and The Logic of Failure by Dietrich Dörner).

In sum, in seeking to avoid failure, the fundamental issue we face is how to manage uncertainty, and not how to calibrate our "risk appetite." Let us now turn to how we naturally react to uncertainty, as individuals and in groups.

Our Complicated Individual Response to Uncertainty

Researchers have found that human beings simultaneously use two separate systems to make sense of the world around them (e.g., see Thinking Fast and Slow by Daniel Kahneman). The first operates quickly and unconsciously, while the second functions more slowly and consciously. Some authors call the first “system 1” or emotion, and the second “system 2” or cognition. System 1 is heavily involved in our individual processing of uncertainty, while System 2 dominates our processing of risk. More specifically, the part of our brain known as the amygdala becomes highly active when we are confronted with uncertainty. The amygdala is a relatively old part of our brain in evolutionary terms, as it developed before the regions that are central to higher rational thought. The amygdala plays a central role in the generation of fear – not just our conscious experience of that emotion, but, more importantly, the range of physiological changes in response to threatening stimuli (e.g., preparations to fight or flee) that precede our conscious recognition of the feeling we call fear or anxiety.

Researchers have found that there are three critical amygdala triggers to which human beings are automatically averse: loss, uncertainty, and social isolation. These operate in a complex manner, not only in isolation, but also in relation to each other. For example, an experience of loss (either material or in terms of social standing) triggers heightened uncertainty and fear of social isolation (e.g., see “Dissociable Systems for Gain and Loss Related Value Predictions and Errors of Prediction in the Human Brain” by Galscher et al). Similarly, heightened uncertainty triggers heightened fear of social isolation (e.g., see “Neural Systems Responding to Degrees of Uncertainty in Human Decision Making” by Hsu et al, and “Neurobiological Correlates of Social Conformity and Independence During Mental Rotation” by Berns et al). From an evolutionary perspective, it is easy to see how these amygdala reactions were highly adaptive, in that they substantially increased the chances our ancestors would survive eons ago on the East African savannah.

In the context of our discussion today, it is also critical to remember that these instincts still operate within us, and that as individuals we have a strong and automatic aversion to uncertainty.

Balanced against this System 1 reaction, however, are a number of System 2 reactions that can often lead us to overcompensate, and excessively reduce our cognitive perception of the amount of uncertainty we face in a given situation. While researchers have sometimes termed these reactions “cognitive biases” and attempted to draw distinctions between them, it has been shown that many of them have common roots in an apparent failure of our brains to give proper weight to bad experiences when updating our assessment of a situation (see “Selectively Altering Belief Formation in the Human Brain” by Shalot et al). This can be seen in a wide range of biases, including “hindsight bias” (where in retrospect we overestimate the probability we attached to events that subsequently did or did not occur); “confirmation bias” (where we give more attention and weight to new information that fits with our existing views, relative to information that does not); the “fundamental attribution error” (where we attribute good outcomes more to our skill than luck, and do the opposite for bad outcomes); the “optimism bias”, where we overestimate the probability of positive outcomes, and underestimate the probability of negative outcomes; and “overconfidence” (believing we are better at something than we actually are).

In a closely related stream of research, John Coates (a former derivatives trader turned Cambridge neuroscience professor) has shown how the short-term reaction of our endocrine system (e.g., cortisol and testosterone) to losses and gains reinforces our “biased updating” process (see his book, [The Hour Between Dog and Wolf -- the Biology of Financial Risk Taking](#)). Critically, Coates shows how once past a “tipping point”, this reinforcement process can work on the downside as well as the upside, and generate feelings of excessive pessimism and depression.

This phenomenon is also explored by Baumeister et al, in their paper “Bad is Stronger than Good”, which finds that, once bad news accumulates to the point that it can no longer be ignored or explained away, its recognition has a much stronger psychological impact than good news. An excellent example of this is recent research into the significant economic and financial market impact of changes in collective uncertainty about economic policy (e.g., “Measuring Economic Policy Uncertainty”, by Baker, Bloom, and Davis, “The Asset Pricing Implications of Government Economic Policy Uncertainty” by Brogaard and Detzel, and “Political Uncertainty and Risk Premia” by Pastor and Veronesi).

Once again, we find that Keynes was well ahead of his time in calling attention to these issues, in Chapter 12 of *The General Theory*: “The state of long-term expectation, upon which our decisions are based, does not solely depend on the most probable forecast we can make. It also depends on the confidence with which we make this forecast – on how highly we rate the likelihood of our best forecast turning out quite wrong. If we expect large changes but are very uncertain as to what precise form these changes will take, then our confidence will be weak. The state of confidence...is a matter to which practical men always pay the closest and most anxious attention... Our decisions to do something positive, the full consequences of which will be drawn out over many days to come, can only be taken as a result of animal spirits – of a spontaneous urge to action rather than inaction, and not as the outcome of a weighted average of quantitative benefits multiplied by quantitative probabilities...If the animal spirits are dimmed and the spontaneous optimism falters, leaving us to depend on nothing but a mathematical expectation, enterprise will fade and die – though fears of loss may have a basis no more reasonable than hopes of profits had before.”

It is also important to note that a number of researchers have highlighted the apparent contradiction between our modern view that rational/optimal thinking involves the accurate incorporation of new good and bad information into our

updated beliefs (technically, this is the view that we should behave as perfect Bayesians), and the fact that our cognitive biases appear to have been hardwired into our brains by evolution, and must therefore in some way be adaptive and advantageous.

A growing body of research finds that this is indeed the case (for broad overviews, see “The Evolution of Misbelief” by McKay and Dennett, and “The Evolution of Overconfidence” by Johnson and Fowler). For example, Johnson, Weidmann, and Cederman found that overconfidence confers an adaptive advantage in war (“Fortune Favors the Bold”). Similarly, Hirshleifer and Luo show how “the long run steady state equilibrium of a financial market always involves overconfident traders as a substantial positive fraction of the population” (“On the Survival of Overconfident Traders in a Competitive Securities Market”). Finally, Goel and Thakor have shown how overconfident managers have a higher probability of being selected as CEO (“Overconfidence, CEO Selection, and Corporate Governance”).

Thus far we have seen how at the individual level we have a strong and automatic aversion to uncertainty, and how various cognitive biases seem to have evolved to enable us to take purposeful action in uncertain situations (though at the cost of systematically underestimating the actual amount of uncertainty we face). We have also seen how, when these cognitive defenses are overcome, powerful swings toward pessimism and inaction can result. Let us now turn to how these individual emotional and cognitive processes change when we are part of a group.

How Group Membership Changes Our Response to Uncertainty

A classic example of how our perspective changes when we are a part of a group involves our feelings of regret and envy. When we are making decisions on our own, research has shown that we tend to be risk averse, because have a strong desire to avoid the feeling of regret that would be caused by a financial loss.

However, when we know that the result of the same decision will be observed by others, we are much less risk averse, because our desire to avoid envy (a relative loss of social position, which would result if another person receives a positive outcome from their decision) is stronger than our desire to avoid regret (a private emotion associated with an experienced financial loss). As Bault et al note in “Interdependent Utilities: How Social Ranking Affects Choice Behavior”, in private settings, anticipated losses loom larger than gains; in social settings, the opposite holds true.

Other research has found that groups are also much less averse to uncertainty than are individuals (see “Comparing Group and Individual Choices Under Risk and Ambiguity”, by Brunette, Cabantous, and Couture). Moreover, this reduction in uncertainty aversion (or increase in tolerance for uncertainty) increases in proportion to a group’s preference for intuitive relative to analytical decision making (see “Manipulating Reliance on Intuition Reduces Risk and Ambiguity Aversion”, by Butler, Guiso, and Jappelli). Uncertainty aversion also weakens as a group becomes more diverse, (see “Ambiguity Aversion, Group Heterogeneity, and Corporate Investment”, by Hilary and Hui).

Rising uncertainty also tends to increase a group’s tendency toward conformity and herding (e.g., see “Decoding Social Influence and the Wisdom of the Crowd in a Financial Trading Network”, by Pan, Altshuler, and Pentland, or “Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market” by Salganik, Dodds, and Watts). We have previously seen how the causes of this phenomenon are deeply rooted in our evolutionary past, and the way System 1 and our amygdala unconsciously raise our fear of social isolation in situations of heightened uncertainty. Various strands of research have explored the different ways this can manifest itself. These include both unconscious channels (e.g., emotional contagion, in the case of mass panics and hysterias, and groupthink), and conscious ones (e.g., maximizing one’s chance of survival, or keeping one’s job, by deliberately not bucking the group’s view, as in Keynes’ famous observation

that “worldly wisdom teaches that it is better for reputation to fail conventionally than to succeed unconventionally”).

These tendencies toward greater group tolerance for uncertainty and increased herding when uncertainty is high can be further reinforced by organizational incentive systems, cultural norms and business practices, including the extent to which discussion and dissent are rewarded or penalized, and whether people with certain characteristics (e.g., a tendency toward optimism, overconfidence, and/or conformity) are systematically recruited and promoted. It is also clear that a CEO can have a very strong impact on these norms (e.g., see “Groupthink: Collective Delusions in Organizations and Markets” by Roland Benabou). In this regard, perhaps the most dangerous situation for a company is when a leader with psychopathic tendencies (among which are extreme insensitivity to uncertainty and risk) is promoted to CEO – something that apparently happens more often than most of us realize (see “The Psychopath in the C-Suite” by Manfred Kets de Vries, and “Corporate Psychopathy: Talking the Walk” by Babiak, Neumann, and Hare).

A final factor that can affect a group’s tolerance for uncertainty is the national culture in which it operates. Geert Hofstede has identified the strength of preference for “Uncertainty Avoidance” as one of the four main dimensions on which cultures differ. Here is how the U.S. Peace Corps training manual describes this phenomenon: “Cultures characterized by high uncertainty avoidance feel especially anxious about the uncertainty in life and try to limit and control it as much as possible. They have more laws, regulations, policies, and procedures and a greater emphasis on obeying them. They also have a strong tendency toward conformity, hence predictability. People take comfort in structure, systems, and expertise—anything that can blunt or even neutralize the impact of the unexpected.” In the most extreme situations, national cultures with high uncertainty avoidance can spawn corporate cultures where there is too little tolerance of uncertainty (and thus too little experimentation and innovation) rather than too much.

In sum, with the exception of countries with cultures high in uncertainty avoidance, groups generally have a higher tolerance for uncertainty than individuals, as well as a dangerous tendency toward increased conformity and herding when uncertainty is high. These tendencies can be further reinforced by organizational incentive systems, and in some cases by the nature of the corporate culture created by a strong CEO.

Broadly speaking, the net result of our individual and group responses to uncertainty is that a corporate director is quite likely to face situations in which the actual degree of uncertainty has been significantly underestimated. Let us now turn to what a director can do to meet this challenge, in order to properly carry out his or her fiduciary duties.

Practical Approaches to Achieving Goals in the Face of Uncertainty

The term “strategy” derives from the ancient Greek word “strategos”, which is often translated as “the general’s art.” In my experience, there are seven critical aspects of the general’s creative art:

- The ability to define a clear purpose for an organization, if one does not already exist. Not a vision filled with platitudes, but rather a simple, clear vision that inspires commitment and great effort, and provides a solid basis for anyone and everyone in the organization to make good decisions and trade-offs.
- The ability to make sense of uncertain situations (that are usually complex and evolving), in order to make assumptions about important trends and potential discontinuities, future success requirements, and emerging threats and opportunities. There are three aspects of this sensemaking competence: the ability to identify the most important elements in a situation, to understand the most important relationships between them (i.e.,

to generate causal explanations), and to predict how a given situation could evolve in response to various actions and future events.

- Given this situation assessment, the ability to define clear, measurable goals (i.e., an end state) that an organization must achieve within a given time frame in order to survive and successfully achieve or pursue its purpose.
- The ability to obtain access to scarce resources, either directly or via alliances, including talent, information, technology, time, and money.
- The ability to design a robust plan that has an acceptable probability of achieving the organization's goals in the face of the uncertainty it faces, given the resources that are, or are reasonably expected to become available.
- The ability to identify and effectively manage the most important uncertainties that could affect the nature or the implementation of this plan, and the achievement of the organization's goals.
- The ability to clearly communicate the previous six elements of a strategy, in a manner that produces a high degree of understanding, commitment, and creativity on the part of everyone who has a stake in its success.

From a "risk appetite" perspective, the essence of the challenge facing a director is (a) to accurately assess the extent of uncertainty that is inherent in a strategy, across its multiple dimensions; (b) to determine if this degree of uncertainty is acceptable to the board, investors, and other stakeholders, given the circumstances; and (c) to decide on the actions to take to maximize the chance of achieving corporate goals in the face of uncertainty.

With respect to uncertainty assessment, from what we know about individual and group behavior, it is best to assume that the degree of uncertainty that is explicitly acknowledged in any strategy presentation is likely understated, perhaps to a considerable extent. This raises the critical question of how to have a productive discussion with management and/or among the board that results in a more realistic and shared understanding of the uncertainties facing a company.

There is a range of techniques I have found to be useful in achieving this objective. The underlying logic of all of them is a systematic attempt to reconcile intuitive and analytical estimates of the degree of uncertainty in a given situation. I will order these techniques from the least to the most complicated and time consuming:

1. **Pre-Mortems.** Following the presentation of a strategy, business, or project plan, ask the management team to assume it is some point in the future, and the plan has completely failed. Have them independently write down a list of reasons why this failure occurred, and what, in retrospect, the team could have done differently to avoid this outcome. Collect and organize the results, then systematically discuss them. This almost always results in better recognition of the sources of uncertainty facing the team, and a better plan for achieving its goals. In essence, this approach takes advantage of our hindsight bias, which causes us to ascribe higher probability and to be able to describe in more detail events which have already occurred than those we see as uncertain.
2. **Testing Against Outside Scenarios.** A number of organizations, such as the U.S. National Intelligence Council and Royal Dutch Shell, regularly publish detailed scenarios for how political, economic, and technological conditions could evolve in the years ahead. Alternatively, the International Futures Model at the University of Denver can be used for free to construct your own global, country, and sector scenarios. A board can ask a management team to predict how a current strategy would fare under each

of these outside scenarios, how it would need to be modified (or replaced) in order to achieve minimum acceptable results, and what indicators need to be monitored to provide early warning of the scenario that is actually developing.

3. **Certainty Equivalent Cash Flow Valuation.** Strategic plans should be accompanied by financial model analyses that assess their expected consequences for future cash flows and valuation. At minimum, these modeling results should show the most likely cash flow over each year of the planning horizon. Often, they will also show an “upside” and a “downside” scenario, based on an alternative set of assumptions (either strategic, financial, or both). After hearing this presentation, tell each board member that they have the opportunity to exchange each projected uncertain annual cash flow for a risk free payment from the U.S. (or other government). Ask them to write down the amount they would accept from the government in exchange for each year’s cash flow (note that these amounts can vary from year to year). Collect and organize these (typically, the risk free payments will be significantly less than the risky payments), and for each year discuss which uncertainties caused board members to accept the lower payment. In my experience, this always surfaces a range of concerns, including both external conditions and internal execution capabilities.

4. **Explicit Monte Carlo Modeling.** In a Monte Carlo model, the assumed values of key independent variables (e.g., revenue growth rate) are expressed not as point estimates, but rather as a distribution of possible outcomes. Some models also use estimated correlation to model the way the values of two or more variables covary with each other. Monte Carlo models express the value of key dependent/outcome variables (e.g., annual free cash flow) not as point estimates, but as probability distributions of possible outcomes. Monte Carlo analysis also highlights the degree to

which variation in different independent variables affects variation in key dependent/outcome variables, and the different combinations of independent/assumption variables that lead to failure scenarios (e.g., running out of liquidity). It is always enlightening to see how far off a strategy's base case assumptions have to be before cash flow problems develop; there is nothing that concentrates the mind of a board and management team than the prospect of the money running out. In my experience, the best board discussions of Monte Carlo analysis results are based on models that include the strategic and operational variables that drive financial results (i.e., they are based on more than just financial variables, such as the probability a competitor takes a certain action, a regulatory change takes place, or an R&D project meets technical, cost, and schedule targets).

- 5. More Complicated Modeling.** For all its strengths, Monte Carlo modeling also has some important weaknesses. For example, it is hard to implement uncertainty about the model itself (i.e., the independent variables it contains and the relationships between them), or its evolution over time (i.e., in the nature of the model, in the distribution of values for independent variables, or changes in the relationships between them) without a step function increase in model complexity (which tends to detract from the quality of board level discussion about the results it produces). Yet in highly uncertain situations, it sometimes makes sense to build a more complex model that does a better job of addressing these uncertainties, not necessarily to identify an optimum solution, but rather to aid in the search for solutions that are robust (i.e., those likely to achieve minimum acceptable results under a wide range of possible future scenarios). A number of alternative approaches have been used to construct such models, including agent based modeling (e.g., see "Complexity Economics" by Brian Arthur), system dynamics modeling (e.g., see systemdynamics.org and John Sterman's work

at MIT), and Robert Lempert's work at the RAND Corporation, (see his article, "A New Decision Sciences for Complex Systems").

Once a director has a better idea of the true extent of the uncertainty that surrounds a proposed strategy, his or her next challenge is to determine if this degree of uncertainty is acceptable. Unfortunately, there is no neat formula for answering this question, and different groups of stakeholders may hold very different views. In my experience, it is often the case that circumstances dictate the answer. In a famous formula credited to Art Lykke of the U.S. Army War College, the intuitive riskiness of a strategy is a function of the degree of perceived imbalance between the goals to be achieved, the resources available, and the plan for employing them. Some commentators have also noted that the degree of uncertainty in the assumptions that underlie the strategy further contributes to its riskiness. Throughout history, commanders' willingness to accept higher degrees of risk has closely correlated with the importance of the goals they seek to achieve; for example, when survival is at stake, willingness to pursue a risky strategy is necessarily high.

This process of "backing into" a risk tolerance decision based on the importance of the goal being pursued also has analogues in our personal lives. For example, most people would not rush into a burning building to save a computer; however, if their child is inside they would not hesitate.

But what about the majority of cases where survival is not at stake? In my experience, the answer here depends on a more complicated set of calculations.

Consider another example from our personal lives, in this case, a typical retirement saving problem. A person knows how much money they make today, and an expectation about how that may change in the future. He or she also has a target retirement date, an idea of how long they are likely to live after retirement, an idea of the amount of annual cash flow their desired post-retirement lifestyle will require

(and how much they might make through part time employment) and, possibly, an idea of the size of the bequest they would like to leave to their heirs, and/or to charitable causes. With all this information in mind, our retirement saver needs to make decisions about how much to save each year, and how to allocate his or her investment portfolio across multiple asset classes. The textbook approach to this problem is for an advisor to give our retirement saver a “risk tolerance” questionnaire, and based on the results determining an asset allocation and recommending changes in other decision variables. In many cases, however, retirement savers aren’t happy with the result – e.g., a downsizing in post-retirement lifestyle, waiting longer to retire, cutting the size of desired bequests, or cutting current spending in order to increase saving. In many cases, retirement savers are willing to accept riskier asset allocations in their portfolio in order to avoid changes in other decision variables. Behind this decision, however, there is usually a more complicated set of considerations, including recognition of the investor’s ability to absorb and adapt to the risk that the chosen asset allocation will produce lower than predicted returns. These include the options to retire later, work more in retirement, downsize post retirement spending, and reduce the size of desired bequests.

These same concerns with resiliency and adaptability are found in corporate settings.

Despite our best efforts at prediction, we know that in a complex adaptive system, our uncertainty assessments are more likely to be wrong than right. That is not to say that they are useless, or that we should not undertake them. However, when deciding whether a given level of uncertainty is acceptable, we need to act in a manner that recognizes that our forecasts are probably inaccurate, and perhaps significantly so. As in the personal example, directors therefore need to base their “risk appetite” decisions on their confidence in their organization’s ability to detect, absorb, and recover from negative surprises.

Early Detection. Analyses of organizational and system failures that emerge from “unknown unknowns”, from surprise attacks to industrial catastrophes, almost always find that weak signals indicating impending failure were present but went unrecognized. One of the hallmarks of “high reliability organizations” is their consistent focus on identifying these weak signals (e.g., slightly anomalous results, near misses, small surprises, and gut feelings that “something isn’t right”), learning about their root causes, and taking early corrective action.

Commitment to Resiliency. High reliability organizations know that Murphy's Law is an inevitable aspect of all plans' execution. They consistently focus on identifying the mission critical aspects of their plans, and ensuring that backups are in place if something goes wrong. They also recognize that an overemphasis on efficiency can breed fragility, and that slack resources provide a necessary cushion that enables an organization to absorb surprises and survive.

Focus on the Health of Adaptive Processes. If resiliency enables an organization to absorb negative shocks without failing, the health of its adaptive processes determine its ability to recover from them and achieve its goals over time in the face of uncertainty. Broadly speaking, a director should focus on five adaptive processes. I will begin with the one that operates over the shortest time horizon, and conclude with the one that operates over the longest time horizon.

The first and most basic level of adaptation involves constantly seeking ways to make better use of existing organizational resources and capabilities. This is the realm of good day-to-day management. The second level of adaptation involves adding new capabilities to an organization that will create new growth options and enable it to better respond to the full range of possible future scenarios. This is the realm of good talent management and corporate development. The third level of adaptation is focused on improving an organization’s ability to learn, in order to continually improve its performance at an accelerating rate. This is the realm of systematic innovation leadership and management.

The fourth level of adaptation involves changing the organization's definition of success, whether that involves establishing new performance metrics (if the current ones are not aligned with evolutionary selection criteria in the competitive environment), or setting higher goals against current metrics, in order to maintain an appropriate level of "adaptive tension."

The fifth level of adaptation is perhaps the least intuitive. It involves structural changes in the density of the network of connections both within an organization and between it and outside entities – or, if you will, the degree of a system's internal and external connectivity. Let me explain why this is important. Any organization can be described as a series of choices about how to configure different variables. In a corporation, these variables (each of which can be further disaggregated) include marketing, production, distribution, human resources, finance, accounting, information systems, and growth policies. The performance of an organization depends on the choices that are made about how to configure these variables, as well as on leadership's willingness and ability to adapt those choices if the performance they generate is unacceptable.

An organization's degree of internal connectivity refers to the extent to which its ability to make a change in any of these variables is constrained because of the impact of this change on other variables (another way to think of this is how many people in an organization have political veto power over a proposed change). All else being equal, in organizations with densely connected internal networks – i.e., a high degree of internal connectivity -- making a change is often hard and takes a long time, even in the face of declining performance.

An organization's degree of external connectivity refers to the density of its network of connections with outside organizations – for example, the number of customers, suppliers, investors, and regulators it deals with over the course of a year, and the

depth and intensity of those relationships. The higher an organization's external connectivity, the more ideas and requests for changes it will receive.

Now let's consider how internal and external connectivity interact. In the absence of a high degree of adaptive tension, an organization with a high degree of internal connectivity, and a low degree of external connectivity will be excessively stable (or, put differently, have an excessive degree of internal order), and will insufficiently respond when its external environment changes and its relative performance worsens. Most often, because they find change so hard, excessively stable organizations are accused of errors of omission rather than errors of commission. A key indicator of excessive stability is how often an organization finds itself surprised by outside data or events.

At the other extreme, organizations with a low degree of internal connectivity and a high degree of external connectivity tend to undertake too many changes at once, unless the level of adaptive tension is very low. As a result, they tend to feel somewhat chaotic, and are often characterized by a high level of internal disorder, frequent errors of commission, and one "fire drill" and "near miss" after another.

Both excessive internal connectivity and excessive external connectivity are recipes for organizational decline and failure. The adaptive leadership challenge is to maintain levels of internal and external connectivity in rough balance, which maximizes organizational creativity, resilience, and chances of survival.

In sum, in the absence of threats to survival, and a director's willingness to pursue a strategy that involves a high degree of uncertainty should be a function of his or her necessarily subjective assessment of an organization's ability to detect, absorb, and recover from negative surprises. The stronger these capabilities, the greater can be an organization's "appetite" or tolerance for uncertainty.

Conclusion

In many ways, much of the discussion around “enterprise risk management” is misplaced, because it is focused on risk – that is, uncertainties whose consequences and probabilities are relatively easy to quantify. Yet companies most often fail because of their inability to manage true uncertainties, and successfully deal with the “doubts that threaten to block or delay action.” As I have described in this white paper, the individual and group psychological processes at work in when we confront uncertainty are quite complicated. They often result in directors being presented with strategic plans that underestimate the degree of uncertainty that could affect their potential success. The good news is that there are steps a director can take to better assess this uncertainty, to better manage it, and to make better fiduciary decisions about the vexing issue of an appropriate “risk appetite” for an organization.

Tom Coyne and his colleagues advise boards and executives on strategic and investment risk management. He began his career as a credit officer at Chase Manhattan Bank. He subsequently spent almost twenty years as a management consultant, specializing in turnarounds and growth. He has also served as the CFO, CRO, and CEO of public and private companies. He is a member of the Professional Risk Managers International Association, and a frequent contributor to investment research publications, helping investors to make sense of uncertainty and reduce their exposure to severe losses. He provided subscribers with advanced warning of both the 2001 technology bubble collapse and the 2008 global financial crisis. He is also a member of the top ranked team in the Intelligence Advanced Research Projects Agency’s multiyear forecasting tournament. Tom can be reached at tcoyne@sachuestadvisors.com